

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

Memo No. 319

December 1974

LOCALIZATION OF FAILURES IN RADIO CIRCUITS
A STUDY IN CAUSAL AND TELEOLOGICAL REASONING

by

Gerald Jay Sussman and Allen L. Brown

ABSTRACT

This paper examines some methodologies for diagnosing correctly designed radio circuits which are failing to perform in the intended way because of some faulty component. Particular emphasis is placed on the utility and necessity of good teleological descriptions in successfully executing the task of isolating failing components.

Work reported herein was conducted at the Artificial Intelligence Laboratory, a Massachusetts Institute of Technology research program supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored by the Office of Naval Research under Contracts N00014-70-A-0362-0005.

1. Introduction

We believe that problem solutions, whether they are computer programs, electronic circuits, or mathematical proofs, are deliberately designed. Both design and diagnosis (i.e. repair) are processes wherein alternate solutions are proposed, evaluated, and debugged. We think the most effective processes are those we call "Problem Solving by Debugging Almost-Right Plans" (PSBDARP) <Fahlman 1973, Fahlman 1974, Sussman 1973, Sussman 1974, Goldstein 1974a, Goldstein 1974b>. We want to understand the important features of the process of design and the relationship of these features to the organization of deliberately designed structures. We feel that the micro-world of radio circuitry is a reasonably constrained, yet interesting, domain in which to examine PSBDARP <Brown 1974, McDermott 1974>. This paper is an exploration of one important part of the problem of debugging such circuits, that of localization of failures.

There are some very obvious features of any product of PSBDARP. Every such system is made up of distinct parts, be they statements of a program, electronic components, or lines of a proof. Each part has a purpose -- there are no accidental parts (although there may be vestigial parts whose purposes are no longer relevant). Most important, the system must have been debuggable; bugs arising in the design must have been locally patchable. This requires that the system not be completely synergistic. There must be specific aggregations of parts -- modules having distinct and somewhat independent functions. These modules may only interact in constrained ways through distinguished interfaces called ports.

Mathematical proofs are segmented into lemmas, each of which can be debugged privately. Programs have subroutines and macros, often hierarchically arranged. Radios are hierarchically modular. They are divided into sections, stages, networks, and atomic components. Animals too have organ systems, organs, tissues, and cells. This is no accident. The only animals which could have evolved are ones in which a minor design change would not have global side effects.

Parts (modules are higher-level parts) must be described. The description of every part in a deliberately designed system has at least two components -- what it is, and what it does in a particular instance of its use -- the intrinsic and extrinsic descriptions respectively, or in the language of Freeman and Newell <Freeman 1971> "structural" and "functional" descriptions. Thus, a 5 μ Fd, 20V

capacitor (intrinsic description) may serve as an interstage coupling capacitor or emitter-bypass capacitor (extrinsic descriptions). A narrow-band, high gain amplifier may serve as an IF amplifier (extrinsic description). An extensible muscular bag lined with a particular kind of mucous membrane, closable at two ports by sphincters (in short, a stomach) may serve as a vessel in which the first stages of the digestion of proteins is performed. The Chinese Remainder Theorem is intrinsically a theorem about modular arithmetic, but when used as a lemma in Godel's Incompleteness Theorem, it is part of a scheme for encoding and decoding WFF's.

This paper describes the design for a program, LOCAL (part of a larger PSBDAPP system), whose purpose is to localize failures in electronic circuits -- that is, to find the smallest (most embedded) module in a circuit which completely contains the failure. The program we shall describe will be able to diagnose a wide range of radio circuits. It is not designed with any one circuit in mind. The wide range of applicability precludes a diagnosis solely by methods of table-look-up (though it is possible to diagnose a particular circuit this way.) Rather, it encourages abstracting diagnoses from an understanding of the circuit's design. Thus we hope to learn how the principles of electronic circuitry relate to the principles of failure localization in deliberate systems.

Modules at different levels of organization of a hierarchical system may require vastly different analytical techniques for thinking about those modules. A complex computer program written in LISP may be described with bindings, functions, conditional expressions, etc. At the level of machine implementation of LISP we see garbage collection, interrupts, and two's complement arithmetic. The machine is made of registers, busses, ports, etc. The logic is made of TTL, CMOS, or transistors. Transistors are understood in terms of statistical mechanics, quantum theory, and Maxwell's Equations. In animals there is a clear jump at the boundary of organ and tissue. Morphological considerations prevail in discussing organs and organ systems; biochemical considerations are dominant at the level of tissue and cell. In radios too, there seem to be at least two distinct domains. Stages and sections are the domain of signal processing. We speak of mixers, oscillators, amplifiers, and detectors as operating on signals. Intra-stage analysis, in contrast, is the domain of voltages, currents, and impedances.

LOCAL is a hierarchal structure of experts, one for each generic class of module. The structure

is locally imposed in the sense that the decision as to what expert should be called next resides largely with individual module experts rather than some external agent. Figure 1.1 illustrates the basic layout of all experts. The general mode of operation is as follows: When LOCAL is presented with a module suspected of malfunctioning in a specific way, the expert for finding bugs in modules of that class is called with a description of the symptom(s) observed at the ports of that module. Thus the RADIO expert may be called with the symptoms NO OUTPUT or DISTORTION ON STRONG SIGNALS. The POWER-SUPPLY expert might be called with the symptom INCORRECT VOLTAGE. Experts are loath to put the blame on their associated modules. Hence when called with any symptom, an expert first checks its module's inputs and then verifies that the claimed symptom is really there. (This determination will be further explained in section 3.) The expert, having convinced himself of the symptom, must then pin the blame on some submodule of his module. Fixing the blame requires proposing a candidate. Several proposal techniques are utilized including *a priori* probabilities of failure, "knowing the answer," matching the complaint against the extrinsic purposes of submodules, and tracing. Given a proposed failure mechanism, LOCAL must check that the failure could lead to the observed misbehavior, i.e. is it a satisfactory explanation? (Satisfaction must be determined with respect to the particular circuit being examined since the proposal may have been based on general principles that are not applicable in the present instance.) This entails forward causal reasoning that may be trivial, as is the case when tracing, or quite complex, as will be evident in some of the intra-stage debugging scenarios that we will see shortly. The final step is to verify that the claimed trouble is the actual trouble. This verification step is a recursion step, for another expert (associated with the newly proposed failing submodule) is invoked with the failure complaint. Recursion terminates on invoking an expert who cannot localize the problem to some more embedded submodule.

An obvious case of this is an expert for an atomic component, e.g. a transistor. A more subtle case -- suggested to us by Marvin Minsky -- is failure due to general overheating wherein all modules are to blame; hence the recursion terminates with the RADIO expert.

Since each step in an expert's processing schema may fail, an expert may return failure messages as well as success messages to its caller: "This module is not failing!" "This module is failing because <submodule> is failing in <description> way." "This module is failing but the failure locus cannot be resolved further." An expert may also complain that its caller is unfair -- "This module is not

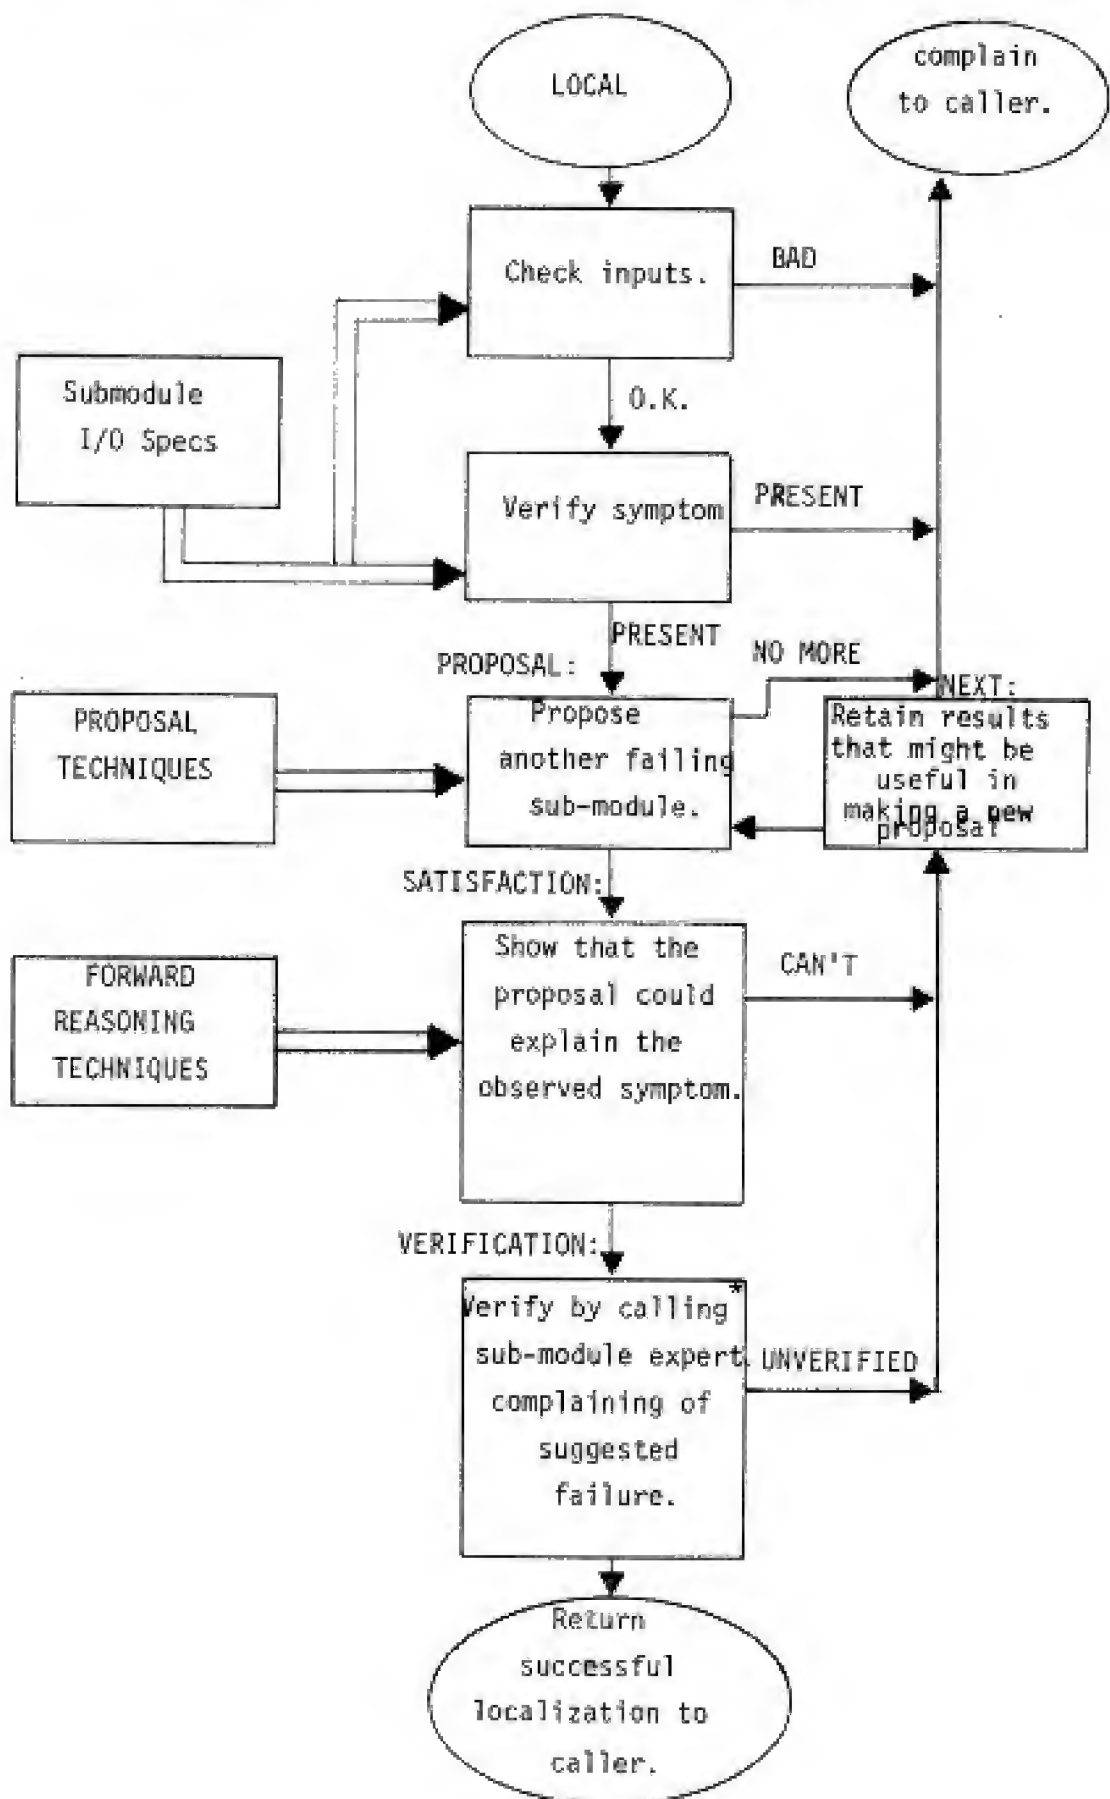


FIGURE 1.1

getting correct input on <port-name>.* We would like to emphasize one additional feature of LOCAL's experts that is revealed by Figure 1.1: an expert factors into 1) an expert-independent control structure that is common to every level of the hierarchy imposed by LOCAL, and 2) declarative and imperative knowledge peculiar to that expert. In the sections that follow we shall be investigating the nature and use of the expert-specific knowledge.

LOCAL will not be able to accept a bare schematic diagram of the radio circuit to be diagnosed. We help LOCAL out by annotating the diagram, parsing it into the module hierarchy implicit in the design represented by the diagram. As part of a PSBDARP system, the plan maker will certainly leave this information on its plans. For reading of externally supplied schematics, we intend to construct a program capable of doing the parsing. The standard names of the modules carry with them various descriptive comments about those modules.

LOCAL communicates to a human assistant (in the same descriptive language as used inside the program) what measurements are to be made, what loops are to be broken and tied off, what signals are to be generated and what parts are to be removed and tested (by applying suitable signals to them). We assume, of course, the availability of whatever signal generators and test instruments are necessary to carrying out the assigned tasks.

2. ANNOTATING CIRCUIT DIAGRAMS

The electronics repairman must be able to understand a circuit in order to repair it. This understanding is reflected in a process of annotating a circuit diagram to indicate the subproblem and solution hierarchy of the circuit designer (see <Goldstein 1974a> and <Ruth 1974> on annotating simple programs). What is the result of this annotation? The purpose of each part in the circuit diagram is determined. An extrinsic comment describing how the part contributes to the proper functions of the next higher level module is attributed to that part. Extrinsic comments are "positive" in the sense that they tell what is to be expected if the design is working properly. Why not also have negative commentary indicating what should happen if a component were to fail? The reason is that the designer is only aware of the positive comments on completing the design. Negative commentary is not (usually) necessary to understanding the design, though there are design features

that prevent bugs from occurring. Moreover, we think that understanding the underlying causes of failures depends on understanding what the circuit ought to be doing.

We do not now know how to write a program that produces a complete annotation from a bare circuit diagram. We do, however, know some of the features such a program must have. As with other processes combining parsing and recognition -- speech recognition and vision for example -- parsing a circuit diagram into annotated submodules is guided at least as much by expectation as by syntax. In a recent "manual" effort at annotating the Heathkit GR-78 receiver <Heathkit 1969>, we were puzzled by the circuit fragment of Figure 2.1. We immediately recognized the module above the

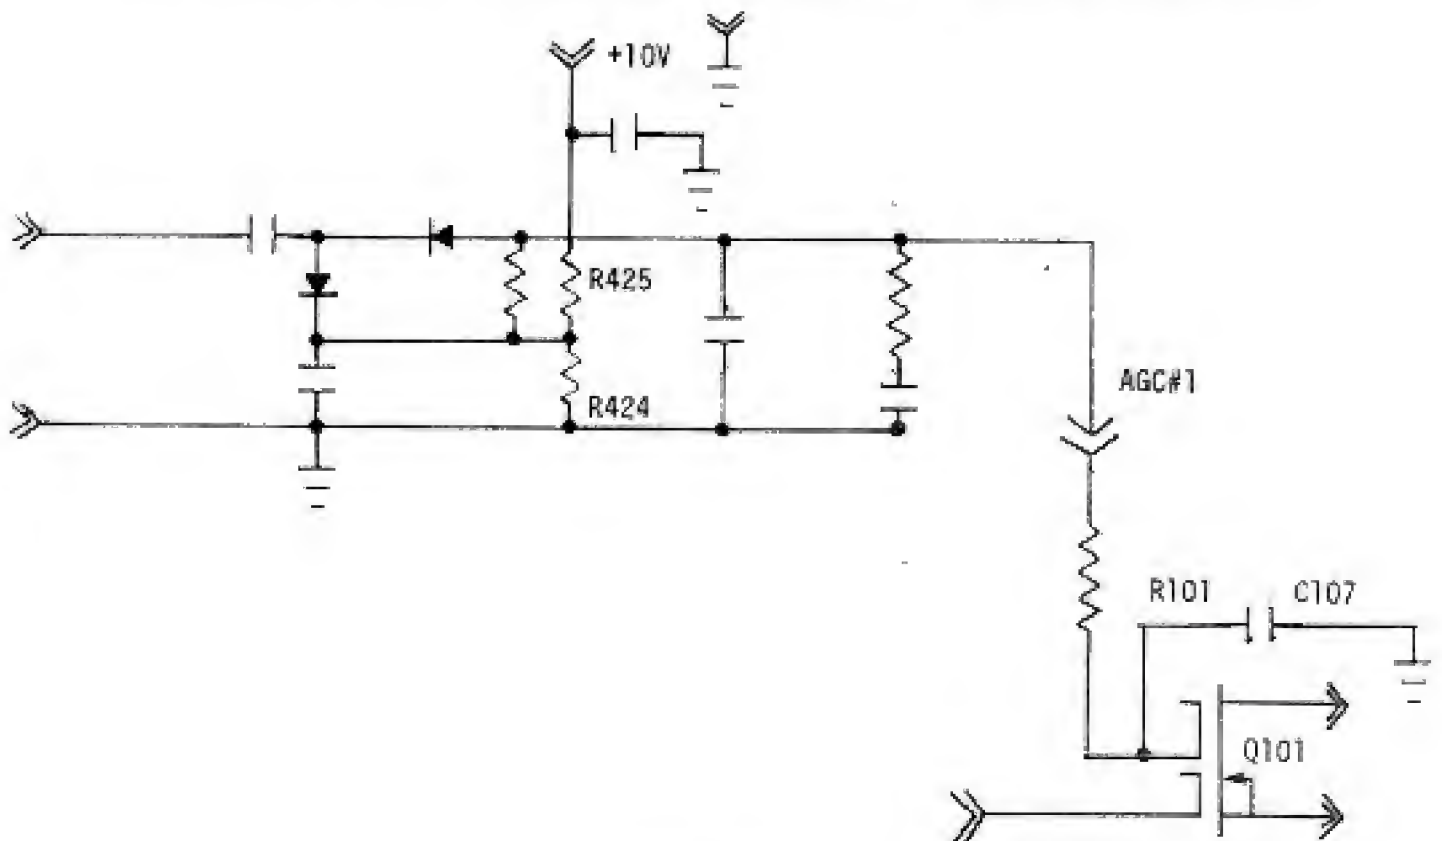


FIGURE 2.1

dotted line as a voltage doubler, except for the resistors R424 and R425. Since the resistors are in series between two constant voltage nodes, it is a good bet that they are a voltage divider. Hence {R424,R425} seems to offset the AC output of the doubler by a DC constant. Why? Eventually we

recalled that we had never found the source of bias for the second gate of the depletion mode IGFET (insulated gate field effect transistor), Q101. (It must have a non-zero bias in order for the IGFET to do anything useful.) The unexplained DC offset could provide such a bias. Since the upper network is an AGC (automatic gain control) and the lower is an RF (radio-frequency) amplifier, our interpretation is indeed correct.

We have in mind a program that could carry out such analyses in the frame paradigm of Minsky <Minsky 1974>. We imagine a voltage doubler frame attempting to explain the voltage offset and an IGFET frame trying to find a bias for its second gate. The two frames meet by virtue of having a common port, and become happy by offering mutually satisfactory explanations; the "unfulfilled expectation" -- bias on the second gate (a constraint) -- matches the "unexplained module" -- the mysterious voltage divider (which must be assigned a purpose). Since a great deal of passive circuitry is concerned with making non-ideal active components look more ideal, a good general strategy for parsing circuit diagrams is to start with the active components and spread outward. A pair of resistors at the base of a transistor and connected to separate DC-fixed nodes is probably a base bias voltage divider. If the transistor, in addition, has emitter resistor, it is probably being operated in a class A regime. Moreover, if there is a capacitor connecting the collector of the transistor to the base of yet another transistor, the capacitor is probably an inter-stage coupling capacitor. This capacitor would define a stage boundary. Understanding circuit diagrams seems to be a matter of parsing top down from the bottom up. By this we mean that we fix on particular components (or collections of components), jump to a conclusion about the use of those components, and attempt to justify the conclusion by locating appropriate additional configurations of components. In justifying a conclusion, additional refining conclusions may be tentatively made. This relaxation <Waltz 1972> proceeds until the tensions caused by unexplained parts are relieved.

In the course of parsing a circuit's components into functional modules, we find that a component (or submodule) may belong to more than one module. A trivial example of this is the inter-stage coupling capacitor. Is it part of the input stage, the output stage, or an entity unto itself? All may be true depending upon the point of view. A similar ambiguity can be seen in the circuit of Figure 2.1 by considering the extrinsic and intrinsic commentary on R101 and C107. Extrinsicly they

form a low pass filter to protect the second gate of Q101 from radio frequency AC. Such AC can arise either from noise on the power input port, or from feedback through Q101 itself. Intrinsically, however, {R101,C107} is a series combination to DC ground from AGC₁'s point of view and a parallel combination to AC ground from Q101's point of view.

We have been assuming until now that the designer supplied circuit diagram has nothing on it but component names and values, and connection information. This is almost never the case. In fact, all of the diagrams we have ever seen indicate the bias values for active components, distinguish the various functional stages, and distinguish various control variables (both internal and external) in the radio. A circuit diagram annotating program can -- and should -- make use of such information. For example, the diagram for the GR-78 indicates that Q101 of Figure 2.1 is part of an RF amplifier. The annotator should know that RF amplifiers typically have variable-tuned circuits at their input and output sides. The complete amplifier, shown in Figure 2.2, has parallel LC combinations at both the gate and drain sides of Q101. The variability of the capacitors clinches the matter and a reasonable annotator would collect those components into modules and comment them as ganged tuning circuits.

Before leaving the matter of parsing, we would like to consider a fragment of a completely parsed radio. Figure 2.3 is a block diagram (a partial functional parsing) of the GR-78. It is a rather sophisticated general coverage, super-heterodyne, AM receiver. This skeleton, to which additional comments will be attached, can be gleaned from explicit commentary found on the manufacturer supplied circuit diagram of the GR-78. (The parser also imposes prejudices of its own in collecting stages into larger modules.) Throughout the rest of this paper, we will be dealing with failures that arise in this particular design. We will expand the stages into detailed subcircuits as it becomes necessary.

We will now consider a detailed annotation of the mixer stage of Figure 2.4. The extrinsic purpose of a mixer is to "mix-down" the modulated broadcast frequency to modulated intermediate frequency. Intrinsically it is a narrow-band (at the output side) RF amplifier whose gain is controlled by the voltage on the oscillator port. {R206,C201,L201} is a narrow-band, parallel tuned circuit, centered at 455kHz. Externically, it provides a narrow-band constriction for the stage's output port. {L203,C206} is a series-tuned circuit centered at 910kHz. Externically it is a wide band frequency

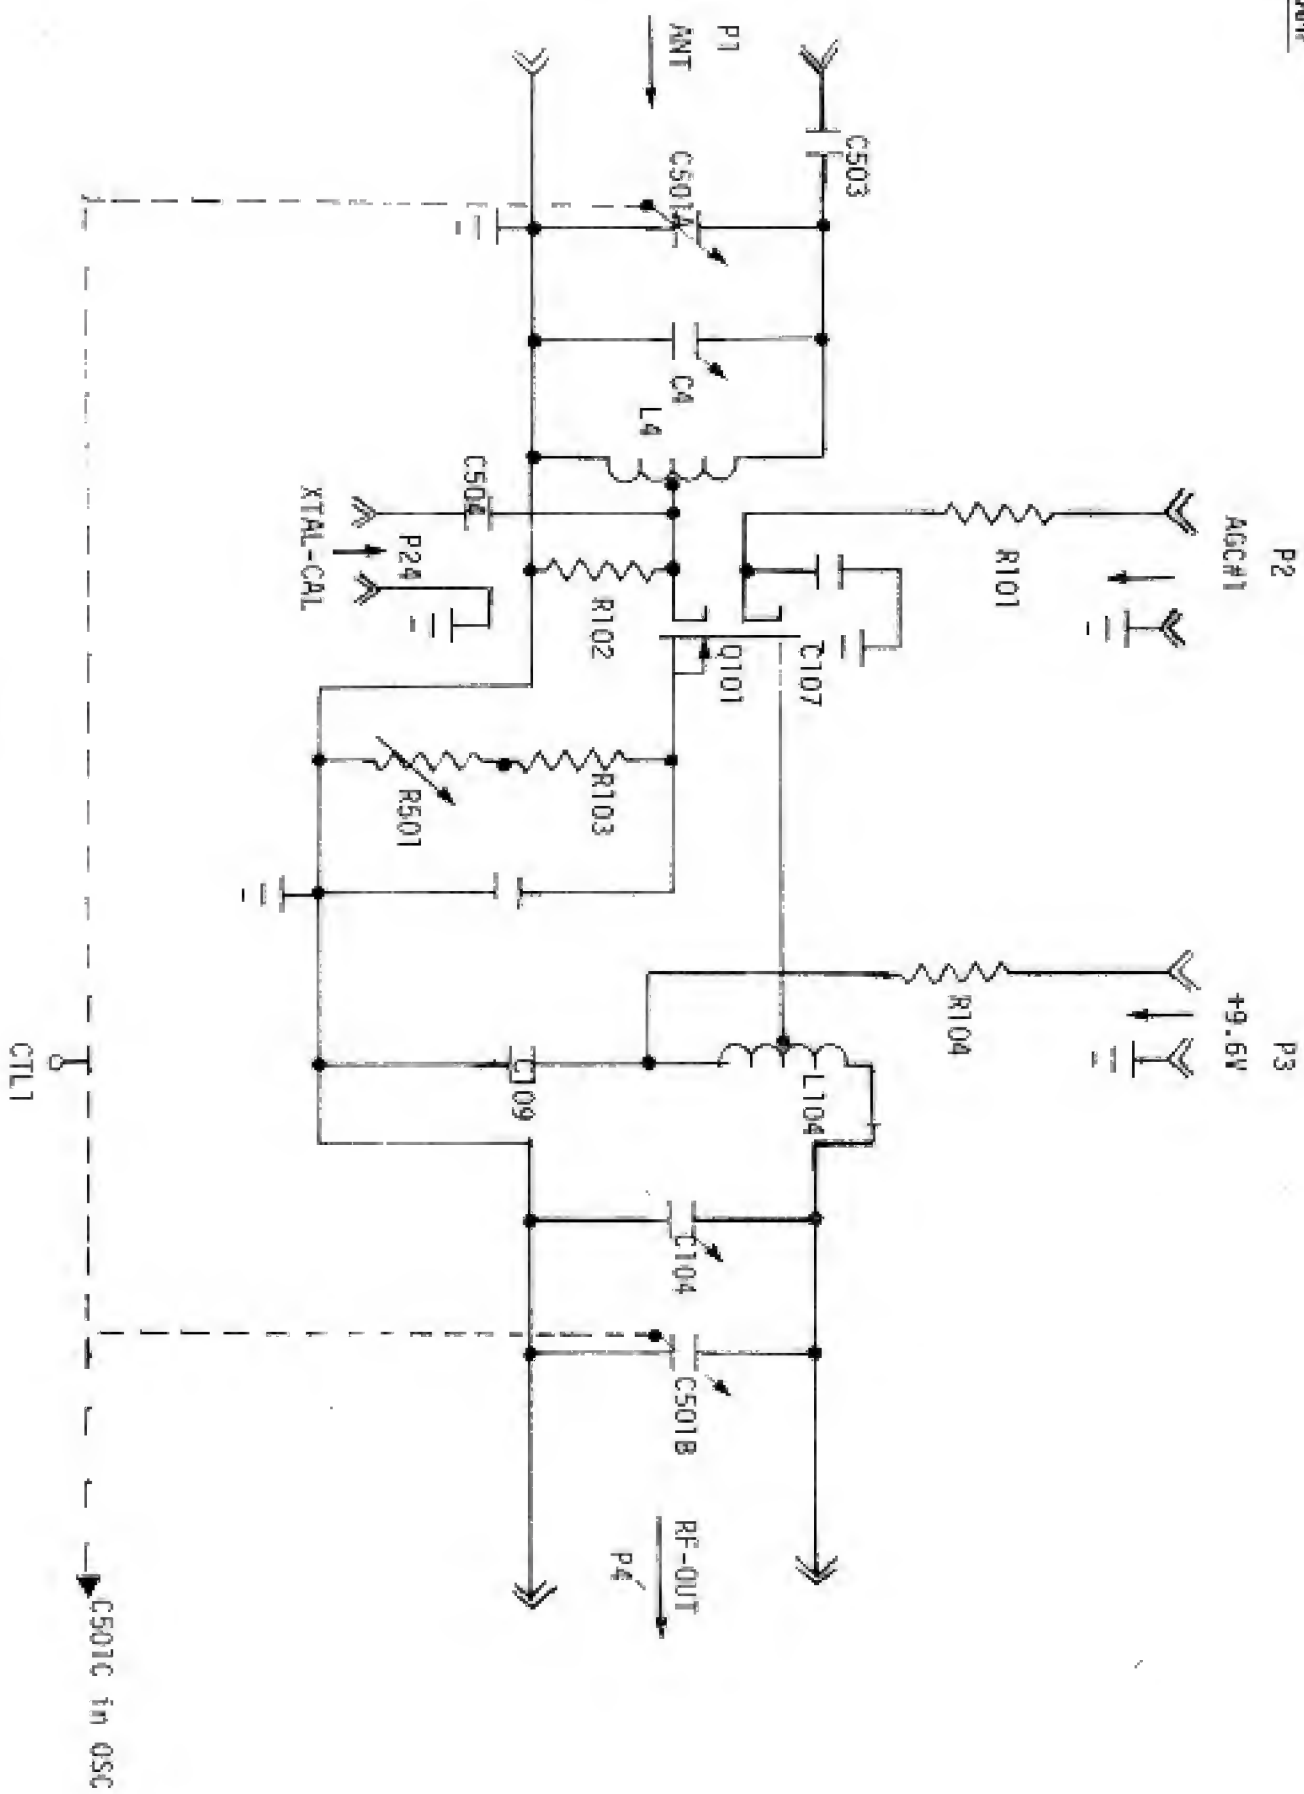


FIGURE 2.2

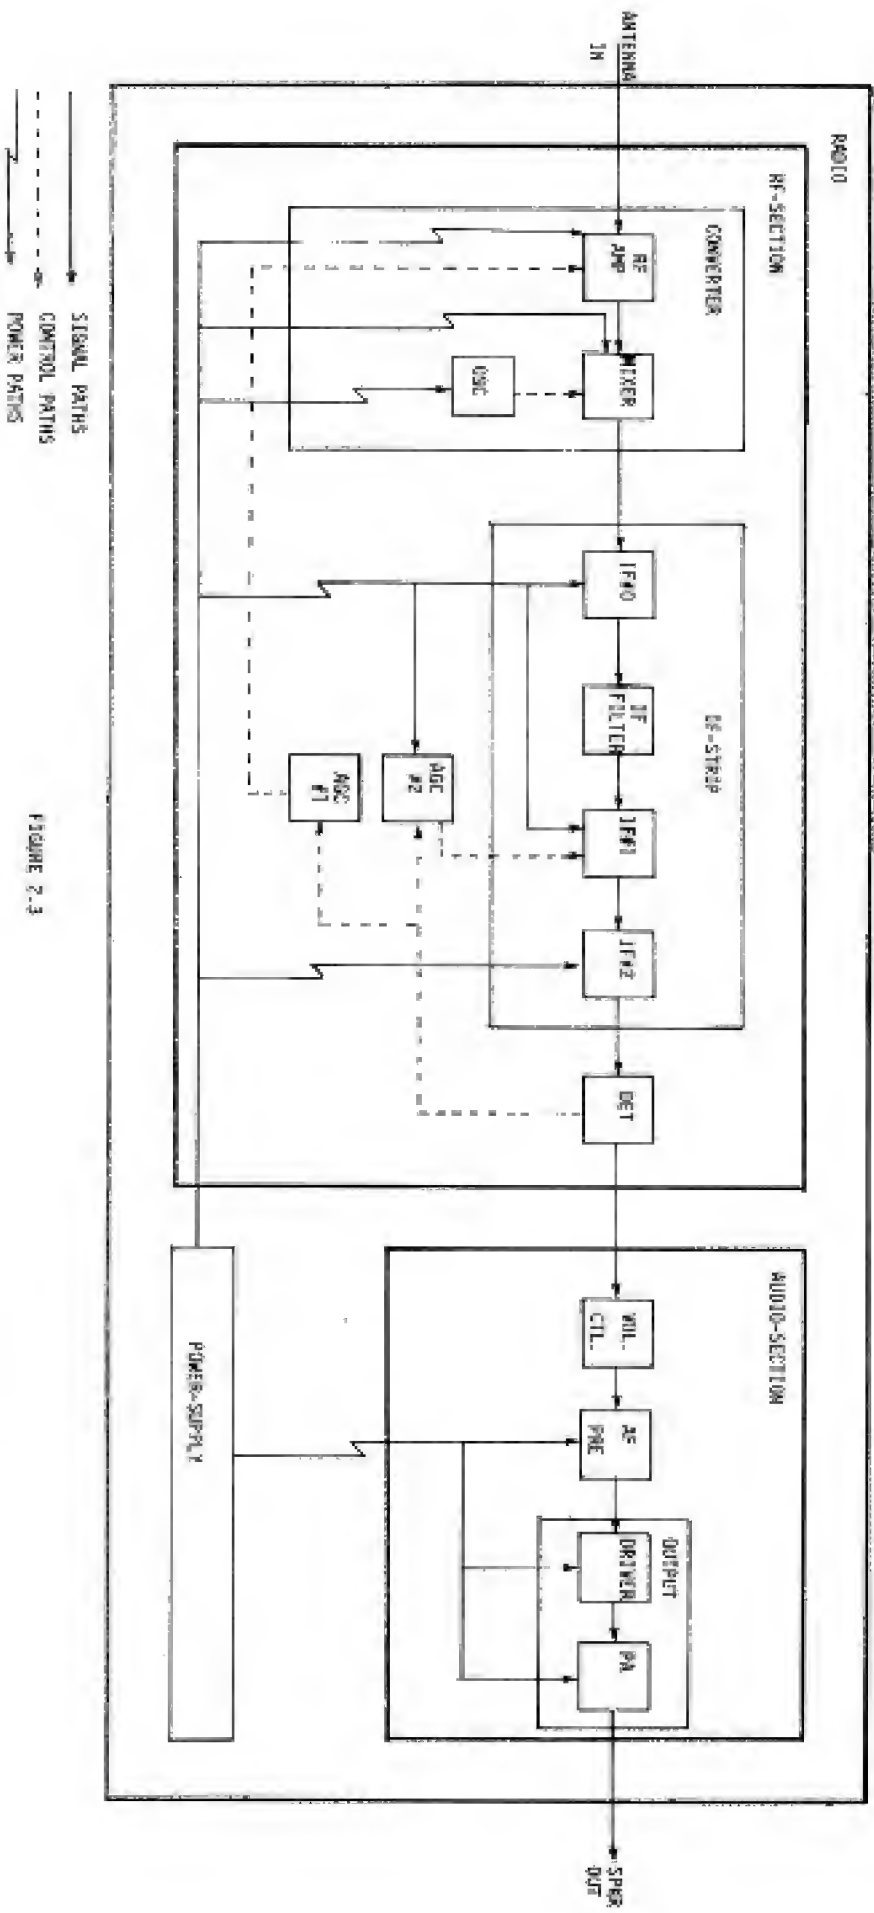


FIGURE 2-3

MIXER

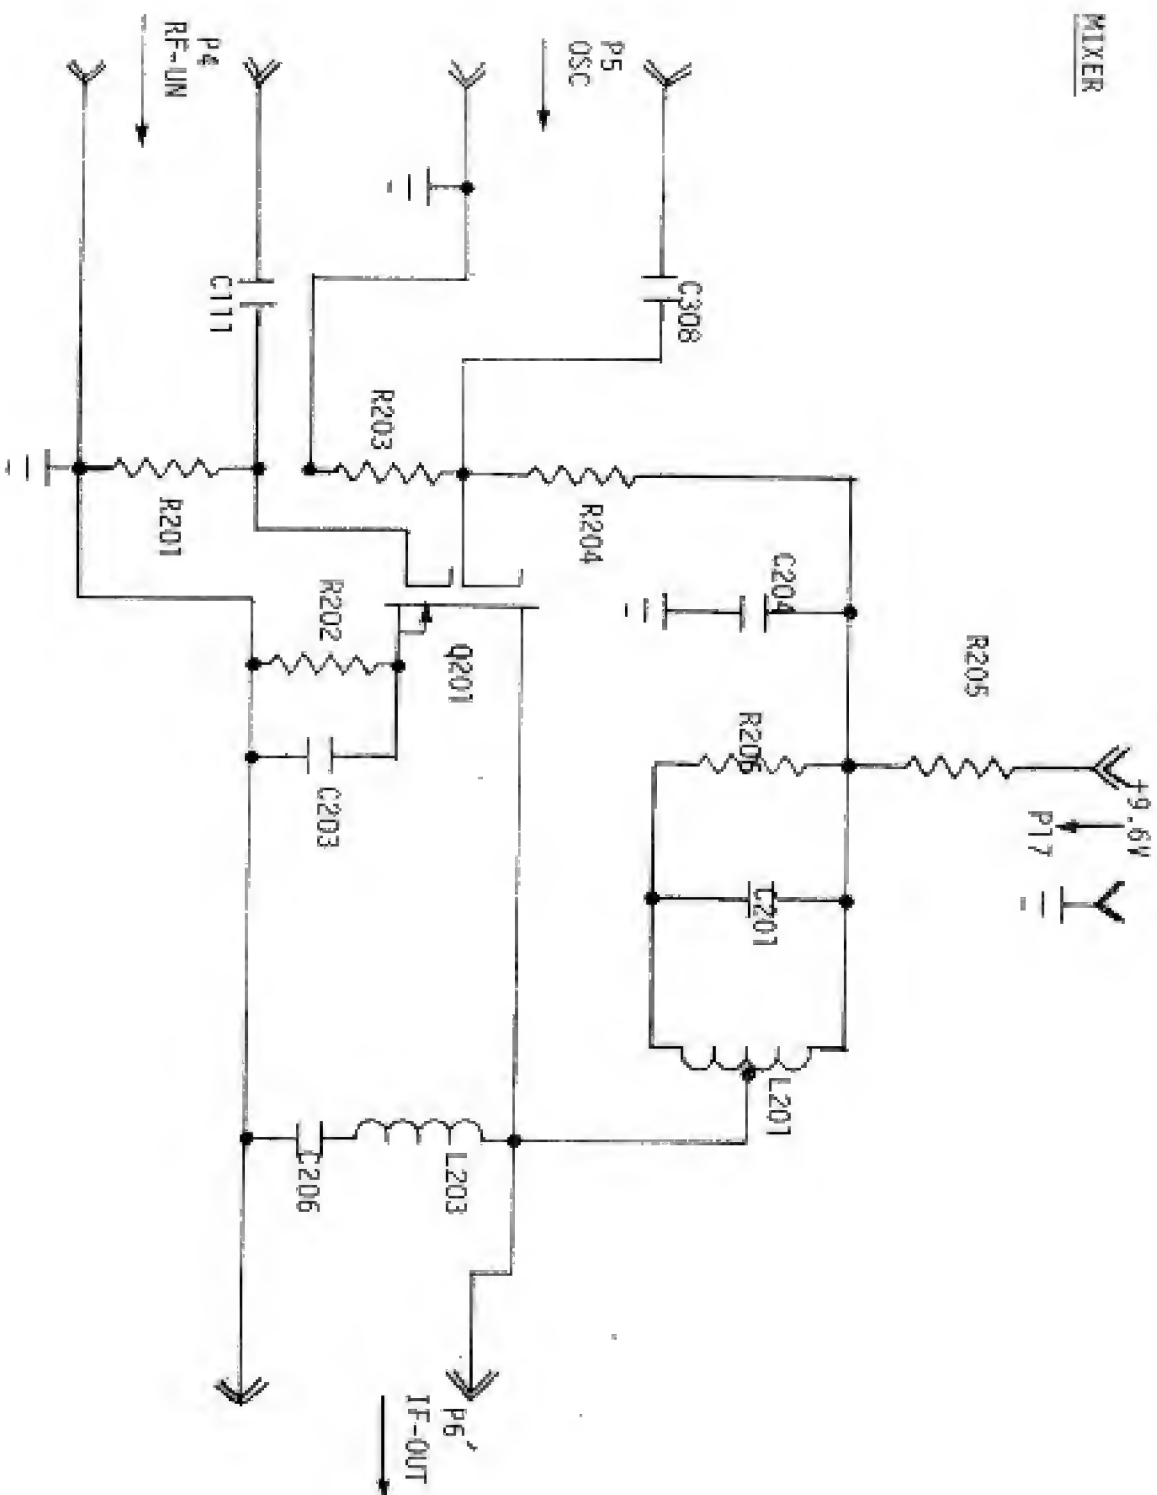


FIGURE 2.4

trap. Q201 provides the stage's variable gain. (We will not give intrinsic commentary for individual components since they are already on the diagram.) R202 is a source feed-back stabilizing resistor. C203 is a source bypass capacitor providing an AC shunt to increase the stage's incremental gain. R201 provides negative bias to the first gate. {R204,R203} is a voltage divider. Extrinsicly it provides the positive bias for Q201's second gate. {R205,C204} is extrinsicly a low pass filter that keeps the top of R204 DC-fixed despite RF on the second gate, the drain, or the power port. From the power port's point of view {R205,C204} is intrinsicly a series combination to ground, while from the oscillator port's point of view it is a parallel combination to incremental ground. C308 and C111 are DC blocking capacitors for their respective ports. One final piece of extrinsic commentary on the mixer is that L201 is center-tapped so as to insure that the Q of the output tuned circuit is set by R206.

3. INTER-STAGE DEBUGGING

Let us examine the localization process in more detail at the level of signal processing -- inter-stage debugging. Suppose the RADIO expert is presented with this radio and the symptom NO OUTPUT. The RADIO expert, as per the scheme of Figure 1.1, checks his power input port (to make sure that the receiver is plugged in!) and then the antenna terminal to see that there is some signal to process. Next he checks the output port and verifies that there is no signal leaving -- as claimed. The RADIO expert cannot open any of his submodules and look into them. He must be content to examine their ports and compare what he sees with what ought to be there.

Because bugs in the power supply have such global consequences, the RADIO expert is prejudiced toward proposing problems in the power supply first. Each power supply output port is proposed in turn as having an incorrect voltage. The proposal is deemed satisfactory if the obvious measurement shows the voltage to be incorrect. (Correctness is deduced from the extrinsic commentary on the power supply.) If any such port does not pass inspection -- say, one to the audio section -- SUSPICIOUS-PORT is applied to it. SUSPICIOUS-PORT mediates between stage experts. It is expert-specific imperative knowledge, e.g. a chunk of code whose use is known to stage experts. As is illustrated below, it catches complaints from inferior experts and recommends to the calling

expert's proposer what port should be suspected next. SUSPICIOUS-PORT determines if either: 1) The fault is in the source of the port (in this case the power supply) for not producing the right stuff *ab initio*. 2) The fault is in a target of the port, say the audio section, for overloading the port. 3) Overloading by the target has caused a failure in the source. 4) Or overdriving by the source has caused a failure in the target. Thus the RADIO expert would first call up (via SUSPICIOUS-PORT) the POWER-SUPPLY expert, then the AUDIO-SECTION expert, as both may have bugs.

If, however, the power supply ports are OK, the RADIO expert proposes that the next place to look is the output module (in this case the audio section). Thus SUSPICIOUS-PORT is applied to the output port of the audio section. Previous verification of the radio's bad output makes an audio section bug a satisfactory proposal. (Notice that since this port has only an input side, SUSPICIOUS-PORT has only the first option listed above.) The AUDIO-SECTION expert, when called, checks his inputs and finds, let us say, no signal on his signal input port. He complains to the RADIO expert (via SUSPICIOUS-PORT), his caller, that he has been unfairly accused. He recommends that SUSPICIOUS-PORT be applied to the audio section's input port -- the port between the RF and audio sections. This causes: 1) the source (RF section) to be suspected of not generating a correct signal, 2) the audio section to be suspected of presenting a low input impedance, thus overloading the port, and 3) (a *priori* unlikely) the audio section has fault 2) causing fault 1). Whether or not the AUDIO-SECTION expert is called to check possibilities 2) or 3) depends upon what the RF-SECTION expert returns. If the latter comes back empty-handed or with problems near his output port (which he will explicitly indicate), 2) and 3) must be investigated. Otherwise, SUSPICIOUS-PORT will be satisfied with explanation 1).

The RF-SECTION expert is called and the bug is similarly traced. Is it clear that this process can be continued recursively until some particular stage (or pair of stages) accepts the blame for the trouble? There is a possible hitch -- feedback. Suppose we come into the RF-SECTION expert with the symptom DISTORTION ON STRONG SIGNALS and suppose that the distortion first appears on the output port of the IF (intermediate frequency) strip. Is it possible that neither the IF strip nor the detector is responsible? Yes! Notice that in addition to a signal port, the IF strip has auxiliary input ports for power and control. A possible cause of the problem is that the automatic gain control (AGC)

buss has become inoperative allowing a strong signal to overdrive IF π into non-linearity. The procedure that we have described can, in fact, catch this. The IF-STRIP expert, on checking his inputs, will discover that the control input is incorrect for the signal input. (Notice that descriptions of control signals must reflect their dependencies on other signals.) The RF-SECTION expert will then pass the trace back; SUSPICIOUS-PORT will be applied to the port between the AGC buss and the IF strip. Now, in fact, the input signal to the AGC buss is distorted. The AGC-BUSS expert, however, does not care what his input signal looks like other than its having some minimal average voltage. Since most any input to the AGC buss will do, the problem must be in the AGC buss or in the IF strip.

Of course, one could imagine feedback situations in which the details of the signal fed back really did matter. In that case some sub-section expert would have passed the trace back to a point along the signal path that had already been visited. This is a tip off to the calling expert that he is in a loop and had better do something to break it. The general solution to the problem is to choose some point in the feedback loop where it can be broken by terminating the output side with the correct load, and simulating a good signal and source impedance at the input side. The decision as to where the loop should be broken will be based on the feasibility of supplying the correct terminal conditions and the mechanical inconvenience of actually executing the break.

How does LOCAL know what signals it should find at the ports of the various modules? This information is implicit in the extrinsic descriptions of the modules. The extrinsic description of the radio is that it is a device that decodes RF encoded audio information into the base audio. This description tells us that the input to the radio must be RF and its output is audio. Now our radio is composed of an RF section whose extrinsic purpose is to sense an amplitude modulated RF voltage and put out an audio voltage that follows the modulation. This latter description not only defines the I/O properties of the RF section but further specifies the input port signal of the overall radio. The RF section has a converter that mixes down amplitude modulated RF to AM-IF at 455kHz. Such characterizations go all the way to the bottom of the hierarchy. The transformational descriptions embodied in the modules' extrinsic descriptions serve to define and refine the nature of the signal at each port. Consequently the obvious first task LOCAL should carry out (on a radio it has never seen before) is to walk over the module hierarchy describing what the signals should be like at the ports.

Note the great power of the extrinsic descriptions: they allow LOCAL to make qualitative predictions about the output signal whenever the input signal is as required. The essential properties of circuits can be predicted more directly than might seem possible (using intrinsic descriptions) by ignoring two kinds of information: the precise and gory details of the signal, and the nature of the output when the input is not as the designer expected. Suppose that LOCAL knows that the mixer stage of Figure 2.4 has at its RF input port a 1000kHz signal, amplitude modulated at 100Hz. LOCAL deduces from the extrinsic description that the signal at the output port is centered at 455kHz and amplitude modulated at 100Hz. Indeed LOCAL could have deduced the same output from the intrinsic descriptions indicating that the mixer is a variable gain RF amplifier with narrow-band output wherein the gain is governed by the sinusoidally varying output of the oscillator. Some algebra and the consideration of the output filter would yield the desired result, but with considerably more work. Of course the intrinsic computation could also tell LOCAL the effect of the oscillator's oscillating at 555kHz as well.

We also think that extrinsic descriptions are the essence of the "structure" in structured programming. Imagine for example a file system having a file deleting subroutine. The fact that in the course of execution this subroutine manipulated certain interlocking mechanisms is irrelevant to understanding its extrinsic behavior as an agent for deleting files.

At this point we must admit that we over-simplified when we explained the intrinsic/extrinsic dichotomy as "what it is" versus "what it does." The deeper one goes into the hierarchy of extrinsic description of the radio, the more one discovers how the radio is implemented. Note also that the tracing technique we have described is applicable to any "flow" processing system where the processing has distinguished stages, e.g. chemical plants, programs, etc.

4. Intra-stage Debugging

After LOCAL localizes the problem to some stage, how does it manage to come up with the offending component? The analytic tools used in inter-stage debugging are not appropriate within stages, as the notions of signal and signal processing give way to the notions of voltages, currents and impedances. Of these three terminal variables only the first is conveniently measurable in an operating circuit. There must be a better way to isolate the failing component(s) than by removing

every component from the stage and verifying its intrinsic specifications.

We expand the proposal/satisfaction/verification processing of Figure 1.1 into the recipe of Figure 4.1. The first order of business, as in inter-stage debugging, is to verify that the inputs are acceptable and that the outputs show the symptoms indicated. The next thing that happens is that the DC node potentials are measured at the internal nodes (nodes not parts of ports) of the suspected stage. (This measuring process may be carried out completely or it may run as a co-routine with the analysis of the hypothesized failure.) If the measured values are substantially the same as the values on the schematic, then more than likely the problem is to be found in some essentially AC subcircuit of the stage. This observation is important to the next step in the recipe, the proposal of a plausible failure in some component. As was the case with the experts encountered in inter-stage debugging, experts at this level have proposers with prejudices about what to try, given various complaints. As a default option, a proposer might offer up failures in their *a priori* order of occurrence. If the DC quiescent values are OK, the proposal algorithm should give higher priority to the consideration of failures of components whose purposes are concerned with the stage's signal coupling to the outside; internal AC feedback paths should be given similar consideration. One might imagine that in the worst case LOCAL blindly verifies the intrinsic specifications of each component. However, most proposals can be eliminated by the analysis packages which will attempt to reason forward from a hypothesized failure to see if it supports the observed AC and DC symptoms of the stage. Finally a component expert is called upon to verify that the indicated component really is the culprit.

This intra-stage localization scheme will work very effectively. When a component fails, the change in its behavior is rarely subtle. Typically a two-terminal component will have shorted or opened. A transistor is likely to suffer the same sort of failure at one or the other of its junctions. A set of such component failures sorted in decreasing likelihood of occurrence might be {transistor, electrolytic capacitor, resistor, ordinary capacitor}. Most component failures lead to significant changes in the quiescent DC conditions of the stage. These changes almost surely engender distortion if not total signal loss. Thus most hypotheses can be filtered out by a rather crude analytic program called BIAS whose behavior we will describe shortly.

Also to be considered are the number of stage experts and their expertise. There are probably

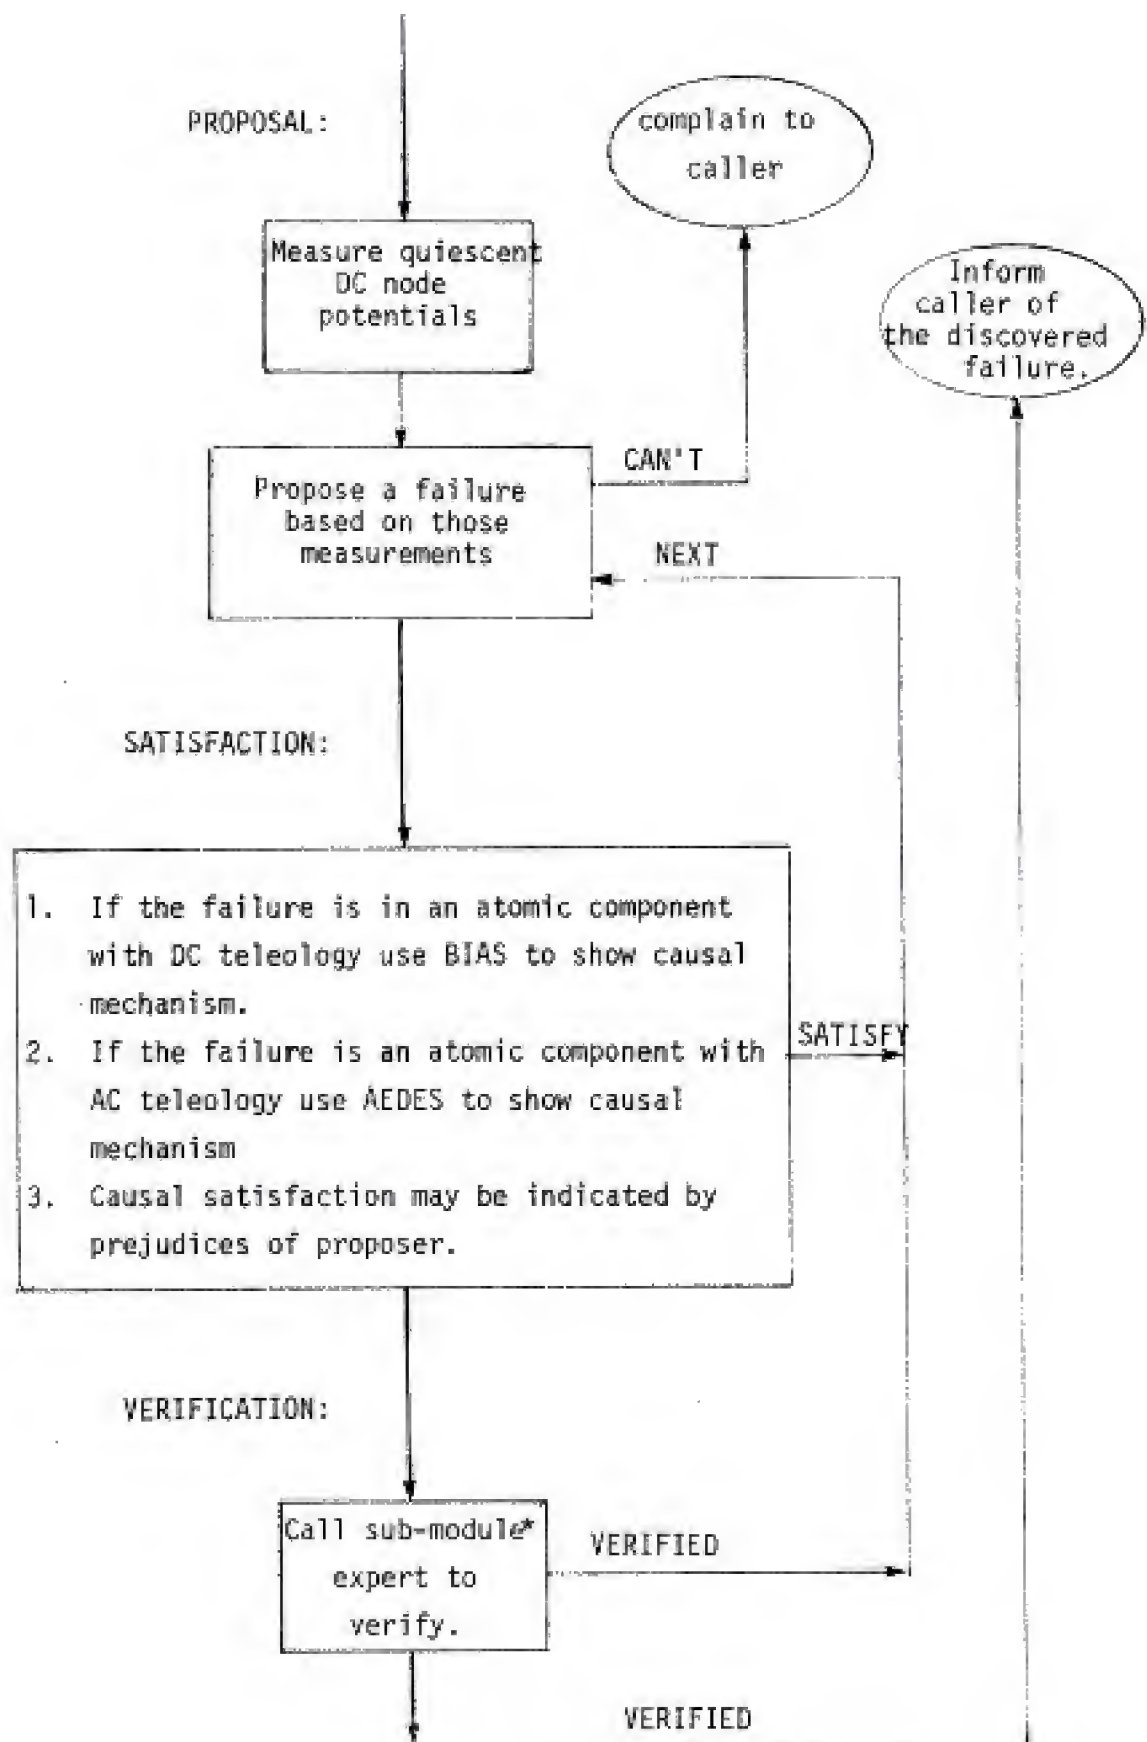


FIGURE 4.1

about two dozen generic stage types to be found in radios. The stage experts' knowledge encompasses the typical manifestations of stage failures, typical stage input and output signals, and typical features of implementations of the stage. Each type of stage can have relatively few kinds of manifestations of internal failures: amplifiers may exhibit distortion. They may oscillate, or they may deliver an output which is uniformly (in the frequency domain) attenuated. An oscillator can deliver a distorted periodic signal, deliver a signal at the wrong frequency (including an aperiodic signal), or deliver no signal at all. A demodulator may exhibit "failure to follow" distortion. As Figure 4.1 indicates, experts are essentially the same from the procedural point of view.

Let us examine the workings of BIAS, which figures most importantly in intra-stage analysis. This analysis is based on the teleology of biasing networks, combined with simple qualitative models of various circuit components.

We know of no presentations of such an informal model in the literature of electronics. Our conversations with people who are active in the field indicate that they all use such models, though perhaps not this one.

We will introduce the circuit model with a mechanical metaphor. It should not be confused with the precise mechanical analogue wherein capacitors, inductors, and resistors are isomorphic to springs, masses, and dashpots respectively. The qualitative model we are using does not explain resonant networks. It does quite effectively explain biasing; the use of reactive components for blocking, coupling, and bypass; and -- to some extent -- the large signal behavior of active components. Consider the common emitter amplifier of Figure 4.2. There are rigid anchor nodes such as $+V_{CC}$ and ground; the stage expert perceives these as being fixed at some DC level. There are floating nodes such as the collector of Q, whose potentials vary incrementally with the prevailing signal conditions. The physical model of a resistor is that when one end of it moves, the other end of it will be pulled in the same direction -- rather like a spring. Inductors and capacitors are similar to resistors, except that their "spring constants" vary with signal frequency. Transistors have a more complex mechanical behavior: a bipolar transistor, operating in the active region (emitter-base junction forward biased and the collector-base junction reverse biased) acts as if the base and emitter were connected by a string about 0.6V long. This means that the base (emitter) may "pull" the emitter (base) but not "push" it. For the purposes of the mechanical model, the connection between the emitter and collector

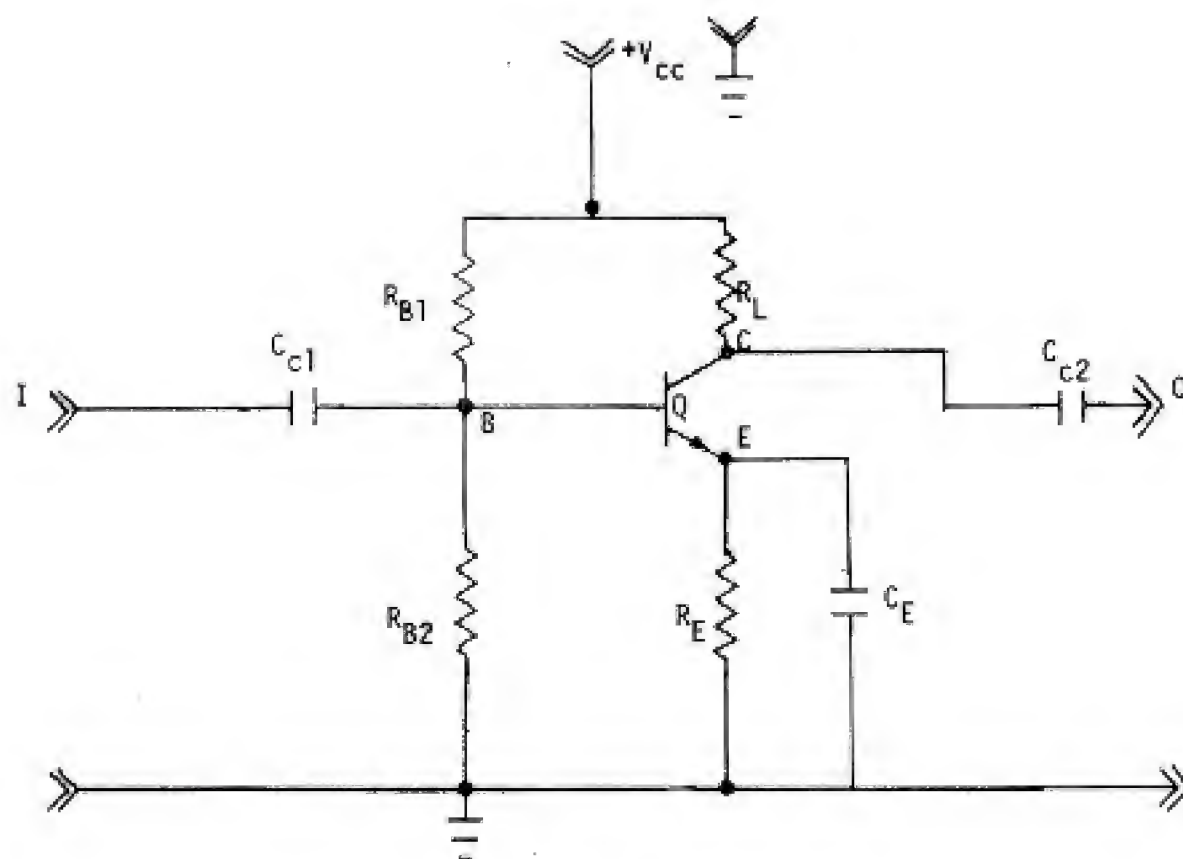


FIGURE 4.2

may be thought of as a resistor (spring) whose resistance (spring constant) varies weakly with base current or emitter-base voltage. In addition to these electronic properties of the transistor, a transistor can be operated in topologically distinct configurations -- common emitter, common collector, and common base -- and in distinct duty cycle regimes -- class A, B, and C. Class A operation is linear, class B operation rectifies but preserves amplitude modulation undistorted, and class C operation throws away all signal information except frequency. BIAS can recognize these regimes from bias considerations.

Now let's see how this qualitative physics is actually implemented in BIAS. BIAS' purpose is to predict the effects of a change at one node on the other nodes of a stage. To that end there are two

important computational structures: the antecedent rule and the propagation path. Antecedent rules describe how a change on one node of a component affects the component's other nodes. Changes are reflected by assertions in a data base. For a resistor the rule is simple: when one node moves (i.e. its node potential changes) the other node moves in the same direction -- but less so -- provided that it is not a static node. BIAS maintains a partial order of the magnitudes of node potential movements which reflects the fact that one node of the resistor moved more than the other. Bypass and coupling capacitors have rules similar to that for a resistor, except that the applied incremental change must be at the stage's associated signal frequency. A more complex component like a transistor must have a collection of rules to reflect this complexity. The data base contains assertions describing the operating configuration of the transistor (e.g. class A, common emitter, etc.) The antecedent rules associated with the transistor, though triggered by assertions about node potentials, must examine the configuration assertions to determine what causal behavior to reflect in the data base. For the transistor of Figure 4.2, a change in the base potential will cause a change of the same magnitude and sign in the emitter potential. The collector potential will move in the opposite direction with a larger magnitude of change.

Propagation paths are constructed by BIAS to make convenient the computation of initial changes in a stage's quiescent DC conditions due to some component failure. Each floating node is examined and the paths attaching it to fixed nodes are noted. (Nodes that are part of the stage's various ports are considered to be fixed from BIAS's point of view.) For example, the base of the transistor has the paths: $\langle B \rightarrow C_1 \rightarrow I \rangle$, $\langle B \rightarrow R_{B1} \rightarrow +V_{CC} \rangle$, $\langle B \rightarrow R_{B2} \rightarrow GND \rangle$, $\langle B \rightarrow Q_1 R_L \rightarrow +V_{CC} \rangle$, $\langle B \rightarrow Q_1 R_E \rightarrow GND \rangle$, and $\langle B \rightarrow Q_1 C_E \rightarrow GND \rangle$. Propagation paths have comments attached to them, indicating, for example, that a path with a coupling capacitor is a DC open unless the capacitor shorts. Similarly a path with a radio frequency choke is an RF open unless the coil shorts. Paths that traverse P-N junctions conduct in only one direction or the other, depending on the prevailing biasing conditions. This means that if a diode becomes reverse biased, its two ends are disconnected in this model.

There is a final rule that should be mentioned before proceeding to examine some analyses carried out by BIAS. Notice that the antecedent rules that we have mentioned hold when a component is true to its intrinsic description in the configuration indicated by the circuit diagram. Not only can

antecedent rules become invalid because of component failures, propagation paths may change as well. Hypothesized failures trigger the manufacture of a new data base context for BIAS, containing different antecedent rules and propagation paths. Antecedent reasoning done subsequent to the hypothesized failure will be reflected by assertions in the new context.

Suppose the presenting symptom given to the stage expert is that the stage (say, an audio pre-amplifier) is exhibiting horrendous distortion. In particular, it seems to be delivering an amplified, but rectified, version of its input signal at its output port.

In our initial presentation of the method of intra-stage analysis we suggested that all the quiescent node potentials may be measured first. In the present illustrations such measurements will be done by co-routine with the analysis and verification procedures. These illustrations are intended to exhibit the proposal of some wrong hypotheses and their subsequent rejection, followed finally by the hypothesis and verification of the actual failure.

The opening of R_L is hypothesized. Hence the node C loses the one path holding it up to $+V_{CC}$. This means that C must fall. Actual measurement of C indicates that it has risen with respect to its nominal value; so this hypothesis is rejected. The next proposal is that R_{B2} is shorted. (This hypothesis would not really come up because of its *a priori* improbability. It is proposed here out of its *a priori* order to illustrate the mechanism.) This means that the node B is shorted to ground. Actual measurement indicates that that has occurred. Since this upsets the forward biasing of the emitter-base junction, the transistor is cut off, i.e. no collector current. Since there is no collector current flowing, there is no voltage drop across R_L , hence C rises to $+V_{CC}$. Actual measurement supports this conclusion. Finally the lack of collector current implies essentially zero emitter current, hence no voltage to speak of across R_E . So the emitter should have fallen to about ground potential. Actual measurement verifies this as well. We now have a hypothesis which is verified by DC forward reasoning but it is unsatisfactory because the shorted resistor couples the AC input signal directly to ground. BIAS's model of the transistor indicates that in the common emitter configuration, the base potential must be incrementally variable in order for the transistor to be operational. If the base is shorted to a fixed point, it obviously cannot vary. So nothing should be on the common emitter amplifier's output port. (We note that there is another component bug that satisfies the DC observations, but not the AC ones: the shorting of the emitter-base junction.) Finally the opening of R_{B1} is proposed. B is held by $\langle B \rightarrow R_{B2} \rightarrow GND \rangle$, $\langle B \rightarrow Q, R_E \rightarrow GND \rangle$, $\langle B \rightarrow R_{B1} \rightarrow +V_{CC} \rangle$ and $\langle B \rightarrow Q, C_E \rightarrow GND \rangle$.

$\langle B \rightarrow R_{B1} \rightarrow +V_{CC} \rangle$ is severed by hypothesis and $\langle B \rightarrow Q, R_L \rightarrow +V_{CC} \rangle$ and $\langle B \rightarrow Q, C_E \rightarrow GND \rangle$ are DC opens by virtue of the collector-base junction and bypass capacitor respectively. Consequently B is held up to $+V_{CC}$ (known to be positive with respect to ground) by one path fewer; hence B must fall. In fact it should fall to ground since it is held by no other positive fixed point. Now the following causal chain ensues: since B has fallen to ground, the base current for Q is zero; hence the collector and emitter currents are about zero as well. Finally there is no voltage drop across R_L or R_E since there is no current in them; so C must rise to about $+V_{CC}$ and E must fall to ground. All these DC conclusions correspond to the DC facts as measured. Given the way the transistor is now biased, only positive swings of the input will result in collector current, i.e. the stage is operating as a class B amplifier. So we have a complete explanation of the original complaint. The resistor expert is called on to verify that R_{B1} has opened. It extracts R_{B1} from the circuit and tests its intrinsic properties. The resistor expert reports that R_{B1} is indeed open, which completes the localization.

The following example shows how a compound failure is handled using the hypothetical context rule we described above. Suppose that LOCAL has just diagnosed an open emitter-base junction and that Q is replaced to effect a repair. The same bug recurs in the "repaired" circuit. LOCAL now must entertain the unpleasant possibility of a compound failure -- one in which the observed failure of Q is a consequence of some deeper cause. LOCAL must come up with such a failure. The failure of Q now becomes a symptom of the as yet unknown cause. LOCAL must enter a new pass of proposal and verification. Suppose in the second pass at proposal the opening of R_{B1} is hypothesized. The reasoning that we saw previously shows that the current flowing in the collector (emitter) would decrease to its operating minimum. This state of affairs is quite unlikely to lead to the opening of a transistor junction that we know to be the case. Hence the opening of R_{B1} is rejected. Next the opening of R_{B2} is proposed. Indeed the base bias potential would rise toward $+V_{CC}$, causing the collector (emitter) current to be large, and having the possible secondary effect of opening the emitter-base junction. Actual DC measurements on the transistor show that the base bias has remained essentially at the level set by the voltage divider. Consequently the opening of R_{B2} must be rejected. Finally the shorting of C_E is hypothesized. The paths holding on to E are $\langle E \rightarrow Q, R_L \rightarrow +V_{CC} \rangle$, $\langle E \rightarrow Q, R_{B1} \rightarrow +V_{CC} \rangle$, $\langle E \rightarrow Q, R_{B2} \rightarrow GND \rangle$, and $\langle E \rightarrow R_E \rightarrow GND \rangle$. Again paths containing healthy capacitors are

ignored. C_E 's shorting means that E is pulled hard to ground. In order for B to track E, the impedance seen by B looking into Q had better be very small. Alternatively, if the voltage drop across R_{B1} becomes big enough to allow B to fall within 0.6V of ground, the current through R_{B1} increases, and the difference goes through the transistor. Thus the collector current of the transistor must increase, increasing the thermal dissipation of the transistor, perhaps exceeding its rating. If no subsequent failure occurred, all the floating node potentials would fall, which is incompatible with observation. Three consequent failures are possible to hypothesize: the opening of the emitter-base junction, the opening of the collector-base junction, or the opening of R_L . Since the first consequence corresponds to the actual state of affairs as analyzed and verified once before, the shorting of C_E is a satisfactory hypothesis. The capacitor expert is called to verify a shorted C_E as proposed.

Let us now consider two other stage types, an IF amplifier and an AM peak detector. The motivation for examining these is two-fold: first we should like to illustrate some fault proposal knowledge that is stage-specific. Also the bugs in these stages will serve as an introduction to a more precise analytic tool, AEDES. Consider the IF amplifier of Figure 4.3. Having localized the failure to this stage, LOCAL presents the stage expert with a complaint of no output at all. Following the recipe presented at the beginning of this section, the stage expert discovers that the bias potentials are all up to specifications (thus the transistor is probably OK). He passes to more specific knowledge about IF amplifiers. He knows that IF amplifiers, being narrow band, usually have narrow band filtering networks at their input and output ports. Furthermore, in order to get a signal out of the amplifier, these filters had better agree as to what band they want to pass. There are two mechanisms that might cause the pass-bands to disagree: the center frequency of one filter has moved substantially with respect to the center frequency of the other. Alternatively, the skirts of one (or both) of the filters may have been "squeezed." Since skirt width is controlled by circuit resistances, the *a priori* probability of the first explanation is much higher. (The foregoing is encoded in the proposal prejudices associated with this stage.) So we need a narrow band filter expert whom the stage expert can ask for failures that might move the center frequency. The filter expert, on receiving the complaint of frequency shift, might try an independent test to verify that that is indeed the case. This could be done by injecting a test signal into the filter. Having decided that

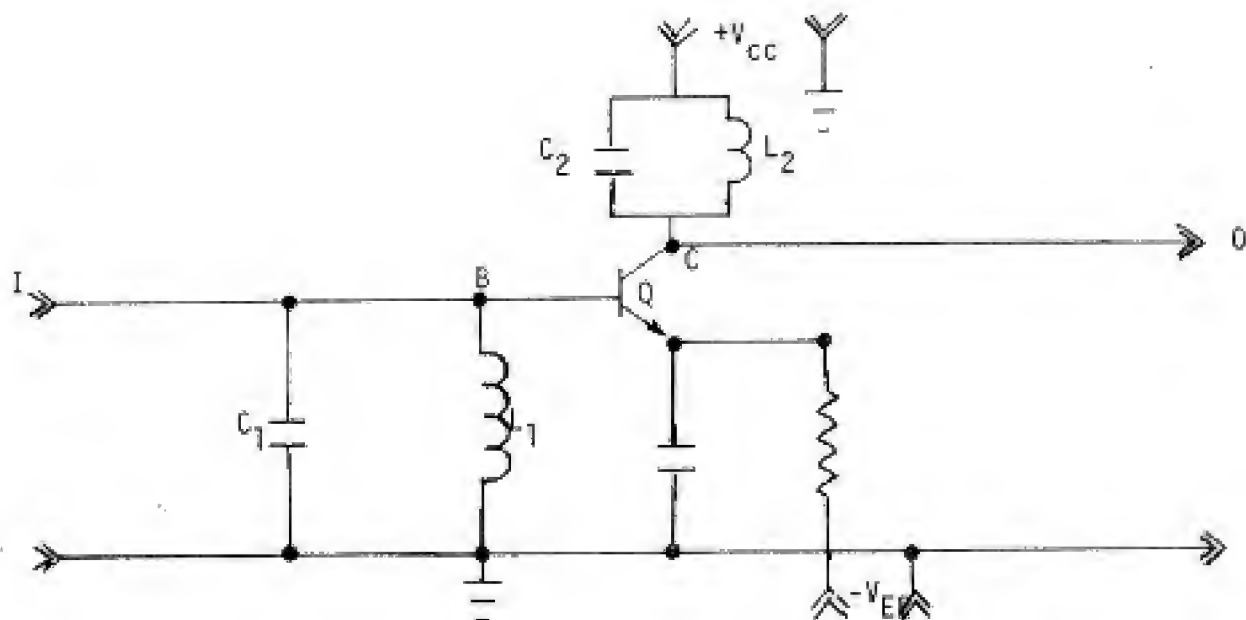


FIGURE 4.3

the center frequency has indeed moved, the precise nature of the failure will be determined with the help of AEDES. The expert does an analysis of the same form as those we have already looked at: proposal, forward reasoning, verification. The difference is that AEDES is the inference tool for forward reasoning. In the present case the failure of some component in the tuned circuit would be proposed. The transfer function of the bug-free tuned network would be computed by AEDES. Then the transfer function of the network with hypothesized bug would be computed. This would enable LOCAL to compare the characteristics of the buggy circuit with the consequences of the hypothesized bug. If they match, the appropriate component expert is invoked for verification.

Now let's look at the narrow-band peak envelope detector stage of Figure 4.4. The operation of this detector (which is reflected in the extrinsic commentary on its parts) can be understood as rectification followed by low pass filtering. Suppose the presenting complaint about the stage were that it had no audio output at all. Indeed some possible underlying causes for this manifestation could be understood in terms of the mechanical metaphor we presented earlier. (By "understand" we mean forward causal reasoning from the hypothesized failure to the observed manifestation.) If the

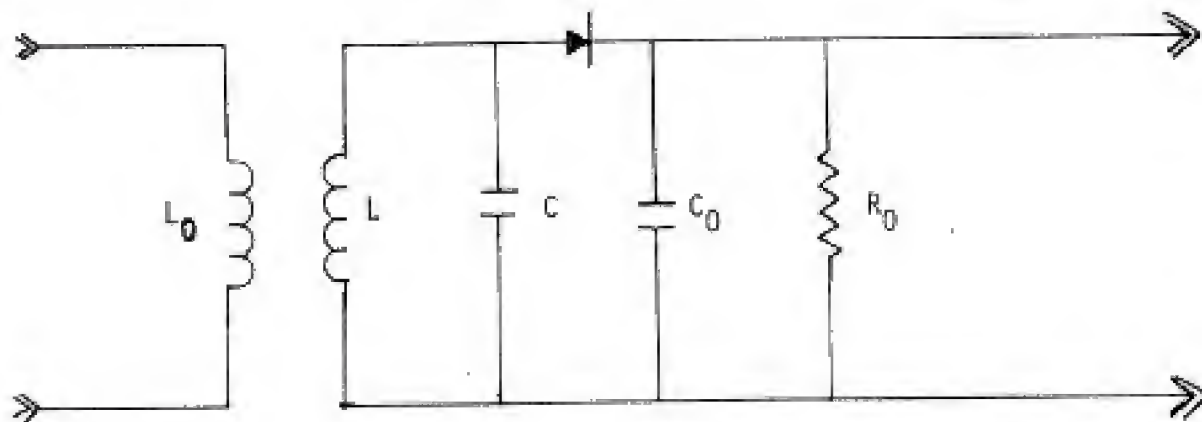


FIGURE 4.4

problem were that the diode had shorted, we doubt that the manifestation could be understood other than by doing a detailed AC circuit analysis. AEDES could be used to compute the transfer function of the detector stage under the conditions of the hypothesized shorted diode. The signal transformation properties of the stage -- with the suggested failure -- could then be compared with the observed transformation properties. If the match is good, the appropriate component expert would be invoked for verification. Notice that this is a particularly trivial detector. A more complex one like Figure 4.5 really calls for rather precise analysis to come to grips with its failure mechanisms, e.g. what happens if one diode opens? We should also point out that there are more subtle bugs in peak detectors than the one mentioned. A detector may exhibit insufficient bandwidth or selectivity. It may also exhibit "ripple" distortion, "failure to follow" distortion, or distortion due to improper offsetting of the diode voltage drop. All of these bugs can be understood only by understanding the detailed AC operation of the circuit.

AEDES is the computational tool for understanding detailed circuit behavior. AEDES models circuits in terms of node equations and system functions. That, however, is where the similarity with other circuit programs ends. (See for example <Penfield 1971> and <Dertouzos 1967>.) AEDES is purely symbolic in two senses. First of all, circuit descriptions are in terms of abstract algebraic parameters rather than numeric descriptions. (Of course the parameters may be bound to numbers and evaluated if desired.) Thus circuit behavior can easily be examined under the variation of parameters, including the extreme conditions of zero and infinity. The second sense in which AEDES is a symbolic

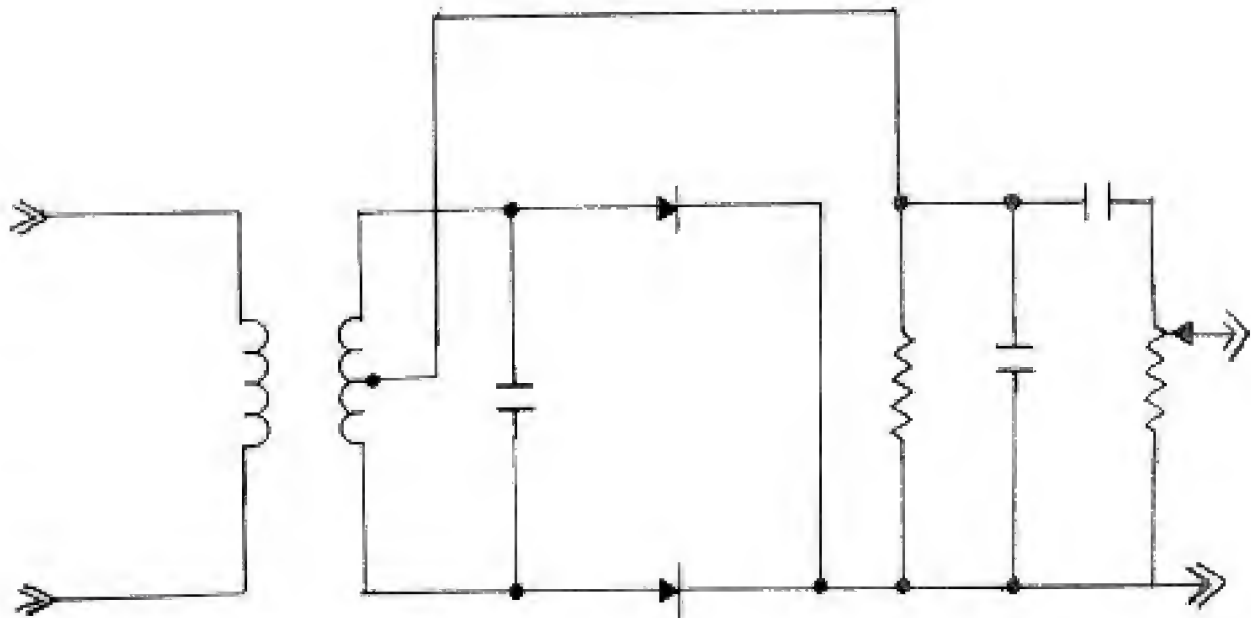


FIGURE 4.5

analysis tool, is that it has access to the hierarchical descriptive structure that LOCAL imposes on the radio. This permits new networks to be analyzed in terms of previously analyzed sub-networks.

5. REFERENCES

<Brown 1974>

A. L. Brown, *Qualitative Knowledge, Causal Reasoning, and the Localization of Failures -- a Proposal for Research*, Working Paper 61, MIT Artificial Intelligence Laboratory, Cambridge, March 1974.

<Dertouzos 1967>

M. L. Dertouzos, *CIRCAL: On-line Circuit Design*, Proc. IEEE, pp. 637-654, 1967.

<Clarke 1971>

K. K. Clarke and D. T. Hess, *Communications Circuits: Analysis and Design*, Addison-Wesley, 1971.

<Fahlman 1973>

S. E. Fahlman, *A Hypothesis-Frame System for Recognition Systems*, Working Paper 57, MIT Artificial Intelligence Laboratory, Cambridge, December 1973.

<Fahlman 1974>

S. E. Fahlman, *A Planning System for Robot Construction Tasks*, Artificial Intelligence, Vol. 5, pp. 1-49, 1974.

<Freeman 1971>

P. Freeman and A. Newell, *A Model for Functional Reasoning in Design*, Proc. Second Intl. Joint Conf. on Artificial Intelligence, pp. 621-640, London, Sept. 1971.

<Goldstein 1974a>

I. P. Goldstein, *Understanding Simple Picture Programs*, Proc. AISB Summer Conference, pp. 37-49, University of Sussex, July 1974.

<Goldstein 1974b>

I. P. Goldstein, *Understanding Simple Picture Programs*, Technical Report 294, MIT Artificial Intelligence Laboratory, Cambridge, 1974.

<Heathkit 1969>

Assembly and Operation of the Heathkit General Coverage Receiver, Heath Company, Benton Harbor, Michigan, 1969.

<McDermott 1974>

D. V. McDermott, *Advice on the Fast-Paced World of Electronics*, Working Paper 71, MIT Artificial Intelligence Laboratory, Cambridge, May 1974.

<Minsky 1974>

M. L. Minsky, *A Framework for Representing Knowledge*, Memo 306, MIT Artificial Intelligence Laboratory, Cambridge, June 1974.

<Penfield 1971>

P. L. Penfield, *MARTHA User's Manual*, MIT Press, 1971.

<Ruth 1974>

G. R. Ruth, *Analysis of Algorithm Implementations*, Technical Report 130, MIT Project MAC, Cambridge, May 1974.

<Sussman 1973>

G. J. Sussman, *A Computational Model of Skill Acquisition*, Technical Report 297, MIT

Artificial Intelligence Laboratory, Cambridge, August 1973.

<Sussman 1974>

G. J. Sussman, *The Virtuous Nature of Bugs*, Proc. AISB Summer Conference, pp. 224-237, University of Sussex, July 1974.

<Waltz 1972>

D. L. Waltz, *Generating Semantic Descriptions from Drawings of Scenes with Shadows*, Technical Report 271, MIT Artificial Intelligence Laboratory, Cambridge, November 1972.

<Watson 1970>

J. Watson, *Semiconductor Circuit Design*, Adam Hilger Ltd., London, 1970.